arXiv:1603.01985v1 [stat.AP] 7 Mar 2016

Multilevel Models with Stochastic Volatility for Repeated Cross-Sections: an Application to tribal Art Prices

Silvia Cagnone *, Simone Giannerini †, and Lucia Modugno ‡

Dipartimento di Scienze Statistiche, University of Bologna, Italy

Abstract

In this paper we introduce a multilevel specification with stochastic volatility for repeated cross-sectional data. Modelling the time dynamics in repeated cross sections requires a suitable adaptation of the multilevel framework where the individuals/items are modelled at the first level whereas the time component appears at the second level. We perform maximum likelihood estimation by means of a nonlinear state space approach combined with Gauss-Legendre quadrature methods to approximate the likelihood function. We apply the model to the first database of tribal art items sold in the most important auction houses worldwide. The model allows to account properly for the heteroscedastic and autocorrelated volatility observed and has superior forecasting performance. Also, it provides valuable information on market trends and on predictability of prices that can be used by art markets stakeholders.

Keywords: multilevel model, hedonic regression model, dependent random effects, stochastic volatility, autoregression.

1 Introduction

The investigation of the relationship between the art market and financial markets has important implications for institutions as well as for auction houses, art merchants and individuals. In fact, also due to the recent financial crisis, there has been a sharp increase in the so called alternative investments that comprise funds specialising in art. These appear to offer a highly beneficial diversification strategy with a complex correlation with traditional assets. Hence, the study of the features of this new asset class is important and cannot disregard the differences with respect to traditional stocks. For instance, art items are exchanged a

*silvia.cagnone@unibo.it
†simone.giannerini@unibo.it
‡lucia.modugno@unibo.it

few times and the transaction costs are considerable. Moreover, there is the so called *aesthetic dividend* which plays a crucial role (see e.g. Goetzmann, 1993; Candela et al., 2013, and references therein).

The study of price determination and price indexes for art items is fundamental to auction houses and art merchants. Besides estimating the price of individual items, price indexes can be used to understand market trends, to assess the main social and economic factors that influence the art market, to understand whether investing in art would diversify risk in a long-term investment portfolio. For art objects, the traditional view of a long run price related to the cost of production does not hold anymore. Some authors argue that the art market is inherently unpredictable since it is dictated by collectors' manias Baumol (1986). However, there is an ever growing consensus, backed by empirical evidence, on the idea that "price fundamentals" can be objectively identified. The hedonic regression, also known as the *grey painting method*, is one of the most used approaches for modelling art prices. It was first proposed in Rosen (1974) and further investigated and applied in Chanel (1995); Ginsburgh and Jeanfils (1995); Agnello and Pierce (1996); Chanel et al. (1996); Locatelli Biey and Zanola (2005); Collins et al. (2009). According to this method, the price of an artwork item depends both on market trends and on a set of characteristics of the item itself. Such dependence is modelled through a fixed effect regression and the estimated regression coefficients can be interpreted as the price of each feature, the so-called *shadow price*. Hence, it is possible to predict the price of a given object by summing the prices of its features. Also, a time-dependent intercept can represent the value of the *grey painting* in that period, that is, the value of an artwork created by a standard artist, through standard techniques, with standard dimensions, etc. (Locatelli Biey and Zanola, 2005). Eventually, the price index is built from the prices of the *grey painting* in different periods.

Despite its potential, the hedonic regression model has several shortcomings. First, as also remarked in Goetzmann et al. (2014) only a small fraction of the great variability of the price dynamics is explained. Second, most of the features are categorical so that the regression equation contains many dummy variables and the resulting models are not parsimonious. Most importantly, the time dynamics is not modelled directly but through dummy variables so that it is not possible to use the model to forecast the prices. Moreover, since it is practically impossible to follow the selling price of each artwork item over time, the available datasets have a structure of a repeated cross-section where at each time point a new sample is observed.

In order to overcome the above-mentioned issues, a multilevel approach can be used (Goldstein, 2010; Skrondal and Rabe-Hesketh, 2004). The treatment of repeated cross-sectional data requires the extension of the classical multilevel model by considering individual heterogeneity within time at the first level, and the variability over time at the second level. This specification has been adopted for the first time by DiPrete and Grusky (1990) and Browne and Goldstein (2010) in the frequentist and Bayesian frameworks, respectively. Motivated by the construction of a price index for auctioned items of tribal artworks, Modugno et al. (2015) extended such multilevel specification as to add an autoregressive

components at the second level. They found a considerable improvement over classical models in terms of prediction and forecasting but the assumption of normality of level-1 residuals was violated, probably due to the presence of heteroscedasticity and kurtosis. They devised an ad hoc solution for deriving robust standard errors through the wild bootstrap scheme for multilevel models introduced in Modugno and Giannerini (2015). Similar findings concerning the non normality and the heteroscedasticity were reported for other specifications for the art market in Bocart and Hafner (2012) and Hodgson and Vorkink (2004). In Bocart and Hafner (2012) the problem was addressed by estimating a semiparametric time-varying volatility and Student's t error with skewness, whereas Hodgson and Vorkink (2004) did not assume any parametric form for the disturbances but retained the assumption of serial independence of random effects. Also, Bocart and Hafner (2013) modeled the volatility of price indexes by means of a smooth function of time as a component of an unbalanced panel model with AR(1) time effects. They implemented a linear Gaussian state-space representation and estimated it through maximum likelihood combined with a Kalman filter, and, as above, they found a violation of the normality assumption. The evidence reported in literature indicates that the volatility of prices plays a pre-eminent role; assuming it constant is not realistic and might cause estimation problems.

In this paper we extend the model proposed in Modugno et al. (2015) by including a stochastic volatility component at the second level by means of a nonlinear state space approach. The specification is motivated by the analysis of the first world database of tribal art prices. This allows to account properly for the heteroscedastic and autocorrelated volatility of level-1 error terms and brings in several advantages. Stochastic volatility models (SV) are based on the assumption that the conditional variance of the observed variable depends on a latent variable that captures the flow of information arriving from the market. Similar to ARCH-type models for financial time series, SV models allow to account properly for fat tailed distributions, white-noise-type dependence, high and lag-decreasing autocorrelations of squared observations. We opt for a stochastic volatility component since ARCH-type models assume that the volatility is affected by past information through a deterministic function. Such a specification is not viable for repeated cross-sections.

Model estimation is performed through maximum likelihood via a non-linear Gaussian filtering process in the spirit of Kitagawa (1987) and Tanizaki and Mariano (1998). The task poses several computational challenges related to the presence of time-varying latent variables that must be integrated out from the likelihood function so that there is no analytical solution. To this aim, Fridman and Harris (1998) proposed a non-linear Kalman filter algorithm by expressing the likelihood function as a nested sequence of one-dimensional integrals approximated by the Gauss Legendre numerical quadrature. Bartolucci and De Luca (2001) extended this approach by computing analytical first and second derivatives of the approximated likelihood. They applied a rectangular quadrature to approximate the integrals. More recently, Cagnone and Bartolucci (2016) approximated such integrals by using an Adaptive Gauss Hermite quadrature

Table 1: Subset of variables classified by type: Physical, Historical and Market.

	Variable	*Categories*
Physical	Type of object	Furniture, Sticks, Masks, Religious objects, Ornaments, Sculptures, Musical instruments, Tools, Clothing, Textiles, Weapons, Jewels
	Material	Ivory, Vegetable fibre, Wood, Metal, Gold, Stone, Precious stone, Terracotta, ceramic, Silver, Textile and hides Seashell, Bone, horn, Not indicated
	Patina	Not indicated, Pejorative, Present, Appreciative
Hystorical	Continent	Africa, America Eurasia, Oceania
	Illustration on the catalogue	Absent, Black/white, Coloured,
	Illustration width	Absent, Miscellaneous, Quarter page, Half page, Full page, More than one, Cover
	Description	Absent, Short visual, Visual, Broad visual, Critical, Broad critical.
	Specialized bibliography	Yes, No
	Comparative bibliography	Yes, No
	Exhibition	Yes, No
	Historicization	Absent, Museum certification, Relevant museum certification, Simple certification
Market	Venue	New York, Paris,
	Auction house	Sotheby's, Christie's

method. Here, we extend the procedure discussed in Fridman and Harris (1998) by using Gauss-Legendre quadrature methods to approximate the integrals involved in the likelihood.

2 The first database of tribal art prices

The first database of tribal art prices was created in 2006 from the joint agreement of the department of Economics of the University of the Italian Switzerland, the Museum of the Extra-European cultures in Lugano, the Museo degli Sguardi in Rimini, and the Faculty of Economics of the University of Bologna, campus of Rimini. For each artwork item, there are 37 variables recorded from the catalogues released before the auctions. The variables include physical, historical and market characteristics. Some of these are shown in Table 1 and most of them are categorical. After the auction, the information on the selling price is added to the record.

Figure 1 shows the boxplots of logged prices aggregated by semester; the number of items sold in each semester is reported inside the boxes. The structure

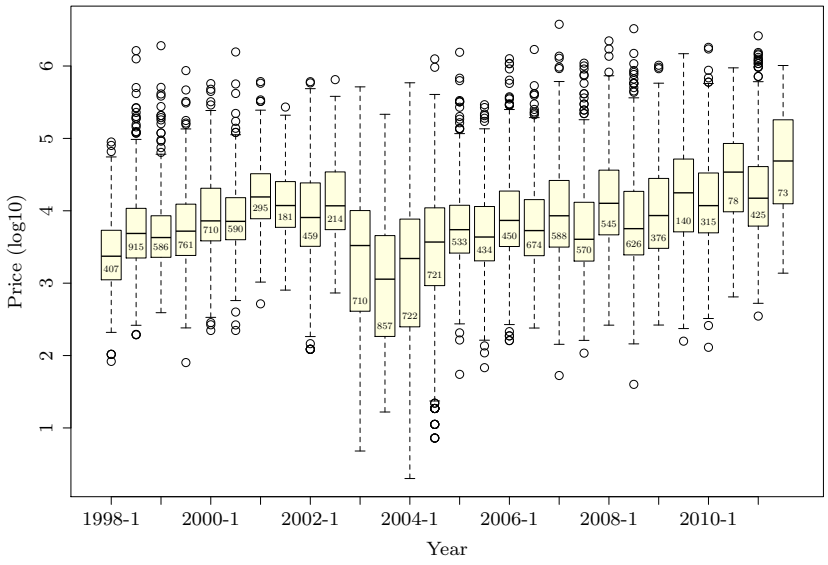

Figure 1: Boxplots of prices in logarithmic scale (base 10) of the tribal art market by semester. The amount of items sold in a given semester is reported inside the boxes.

of the dataset emerges clearly from the graph: in every semester a different group of artworks is sold; e.g. 407 items were auctioned in 1998-1, 915 objects different from the first set were sold in 1998-2, and so on. Hence, tribal art data has a structure like that of repeated cross-sectional surveys and the medians (black lines) give an idea of the trend of prices over time. In particular, note the consistent reduction in the number of auctioned items starting from 2009. Despite this, the overall turnover did not drop since the average price rose. This might indicate the adoption by market agents of hedging strategies against the economic crisis. Overall, we have $T = 28$ semesters, and n_t, the number of items sold in the semester t, varies between 73 (2011-2) and 915 (1998-2); the total sample size of sold items ($n = \sum_{t=1}^{T} n_t$) is 13955. It is convenient to aggregate the data in semesters rather than auction dates since auction dates are not equally spaced in time and this feature is important for modelling time dependence. Also, auctions are organized in two sessions, one in the winter and one in the summer, and each session contains two to four auctions quite close in time. The aggregation in semesters respects naturally this organization so that the results are meaningful from the Economics point of view and this source of heterogeneity can be explained.

3 The model

In Section 3.1 we briefly review the model proposed in Modugno et al. (2015) while in Section 3.2 we extend it by introducing the mixed effects model with stochastic volatility.

3.1 A Multilevel model with autoregressive random effects

Let y_{it} be the observed price for item $i = 1, \dots, n_t$ at time-point $t = 1, \dots, T$ and let $\mathbf{x}_{it}$ be a corresponding column vector of k covariates. Since tribal art data can be thought to have a two-level structure where items represent level-1 units, and time points represent level-2 units, we consider the following random intercept model:

$$\log_{10}(y_{it}) = \beta_0 + u_t + \mathbf{x}'_{it}\boldsymbol{\beta} + \epsilon_{it}, \quad \epsilon_{it}|\mathbf{x}_t \sim \text{NID}(0, \sigma^2)$$

where u_t are time-specific random intercepts that account for the unobserved heterogeneity between items within each time point, $\boldsymbol{\beta}$ is a vector of fixed slopes and β_0 is the overall mean. In repeated cross-sectional data, y_{it} and $y_{i(t+1)}$ are not the price of the same item i observed at successive time points since the two objects are physically different. Conditionally on the vector of covariates $\mathbf{x}_t$, level-1 errors ϵ_{it} (the error term for a given individual at a given time point) follows a normal distribution with constant variance. In other words, the art market is assumed to have a constant volatility over time. Note that ϵ_{it} is conditioned on the vector of the covariates $\mathbf{x}_t$ for all the individuals, that is, we assume strict exogeneity on the explanatory variables (Wooldridge, 2002). Differently from panel data, in repeated cross-sectional data the strict exogeneity assumption implies that, for each item, the covariates are uncorrelated with the error terms.

The dynamics of u_t can be modelled at the second level by extending the above multilevel models as follows

$$u_t = \rho u_{t-1} + \eta_t, \qquad \eta_t|\mathbf{x}_t \sim \text{NID}(0, \sigma^2_\eta),$$

where $\eta_t \perp u_s$ and $\eta_t \perp \epsilon_{it}$ for all $s < t$ and for all i. In this specification, the random effects follow an autoregressive process of order 1 with $|\rho| < 1$, that guarantees stationarity. We denote this model as ARE (Autoregressive Random Effects). Modugno et al. (2015) introduced a full maximum likelihood estimation method via the EM algorithm to fit the ARE model to the tribal art data. The ARE model improves considerably over classical models in terms of prediction and forecasting. However, as it will be shown in the application Section, the assumption of normality of level-1 residuals is violated, probably due to the presence of heteroscedasticity and kurtosis. As mentioned above, the assumption of a constant volatility of prices of art assets is not realistic and might cause severe inference problems. In the following we extend the ARE model by including a stochastic volatility component that accounts properly for the heteroscedastic and autocorrelated volatility of the level-1 error process.

3.2 A Multilevel model with autoregressive random effects and stochastic volatility

We include a stochastic volatility component at level-1 of the ARE model as follows:

$$\begin{aligned} \log_{10}(y_{it}) &= \beta_0 + u_t + \mathbf{x}'_{it}\boldsymbol{\beta} + \exp(h_t/2)\epsilon_{it} && (1)\\ u_t &= \rho u_{t-1} + \eta_t && (2)\\ h_t &= \alpha + \delta h_{t-1} + \sigma_\nu \nu_t, && (3) \end{aligned}$$

for $i = 1, \ldots, n_t$ and $t = 1, \ldots, T$. As before, u_t is the time dependent random effect whereas h_t is the latent variable that represents the volatility component at time t. Both u_t and h_t follow a stationary autoregressive process, so that $|\rho| < 1$ and $|\delta| < 1$. We take the corresponding unconditional distributions as initial conditions for the two processes, namely $h_1 \sim N\big(\alpha/(1-\delta), \sigma^2_\nu/(1-\delta^2)\big)$ and $u_1 \sim N\big(0, \sigma^2_\eta/(1-\rho^2)\big)$. Moreover, we assume that ϵ_{it}, η_t and ν_t are mutually independent, with $\epsilon_{it} \sim \text{NID}(0,1)$, $\eta_t \sim \text{NID}(0, \sigma^2_\eta)$, and $\nu_t \sim \text{NID}(0,1)$ respectively, where, for simplicity of notation, we omit the conditioning upon the covariates $\mathbf{x}_t$. Under these assumptions, the conditional densities result

$$\begin{aligned} y_{it}|u_t, h_t &\sim \text{NID}(\beta_0 + u_t + \mathbf{x}'_{it}\boldsymbol{\beta}, \exp(h_t))\\ u_t|u_{t-1} &\sim \text{NID}(\rho u_{t-1}, \sigma^2_\eta)\\ h_t|h_{t-1} &\sim \text{NID}(\alpha + \delta h_{t-1}, \sigma^2_\nu). \end{aligned}$$

We call the new specification SVARE (Stochastic Volatility and Autoregressive Random Effects) model. The system (1)-(3) is a nonlinear state space representation. Hence, model estimation can be performed by using maximum likelihood via a non-Gaussian filtering process and poses several non-trivial challenges that we describe and address in the following section.

4 Model estimation

4.1 The likelihood function

The maximum likelihood estimation is based on the maximization of the following likelihood function

$$\begin{aligned} L(\boldsymbol{\theta}|\mathbf{y}) &= \int_{\mathbf{h}}\int_{\mathbf{u}} f(\mathbf{y}|\mathbf{u},\mathbf{h}) f(\mathbf{u},\mathbf{h}) \mathrm{d}\mathbf{u}\mathrm{d}\mathbf{h} = \\ &= \int_{-\infty}^{+\infty} \cdots \int_{-\infty}^{+\infty} \left[\prod_{t=1}^{T} f(\mathbf{y}_t|u_t,h_t) f(u_t|u_{t-1}) f(h_t|h_{t-1})\right] \mathrm{d}u_T \ldots \mathrm{d}u_1 \mathrm{d}h_T \ldots \mathrm{d}h_1 \end{aligned} \tag{4}$$

where $\boldsymbol{\theta} = \{\beta_0, \boldsymbol{\beta}', \rho, \sigma_\eta, \alpha, \delta, \sigma_\nu\}$ is the vector of parameters, $f(\mathbf{y}_t|u_t,h_t) = \prod_{i=1}^{n_t} f(y_{it}|u_t,h_t)$ for $t = 1, \ldots, T$, $f(u_1|u_0) = f(u_1)$ and $f(h_1|h_0) = f(h_1)$.

The computation of $L(\boldsymbol{\theta}|\mathbf{y})$ requires solving a $2T$-dimensional integral which is computationally unfeasible. We address the issue by applying an iterated numerical integration procedure introduced by Kitagawa (1987) for non-Gaussian filtering problems. The procedure is based upon rephrasing the likelihood (4) as

$$L(\boldsymbol{\theta}|\mathbf{y}) = \int_{-\infty}^{+\infty}\int_{-\infty}^{+\infty} f(\mathbf{y}_1|u_1,h_1)f(u_1)f(h_1)\int_{-\infty}^{+\infty}\int_{-\infty}^{+\infty} f(\mathbf{y}_2|u_2,h_2)f(u_2)f(h_2)\dots$$
$$\dots\int_{-\infty}^{+\infty}\int_{-\infty}^{+\infty} f(\mathbf{y}_T|u_T,h_T)f(u_T)f(h_T)\mathrm{d}u_T\mathrm{d}h_T\dots\mathrm{d}u_2\mathrm{d}h_2\mathrm{d}u_1\mathrm{d}h_1. \tag{5}$$

The resulting bivariate integrals can be approximated by using numerical quadrature techniques. The most common quadrature techniques for stochastic volatility models are the rectangular quadrature (Bartolucci and De Luca, 2001) and the Gauss-Legendre quadrature rule (Fridman and Harris, 1998). We adopt the latter method since in this context we found it superior with respect to other proposals, especially in terms of computational time. The application of the Gauss-Legendre quadrature rule to eq. (5) produces the following approximated likelihood function

$$\begin{aligned}\tilde{L}(\boldsymbol{\theta}|\mathbf{y}) =& \left(\frac{b-a}{2}\right)^T\left(\frac{e-d}{2}\right)^T\sum_{i_1}^{n_u} w_{ui_1}\sum_{j_1}^{n_h} w_{hj_1} f(\mathbf{y}_1|u^*_{i_1},h^*_{j_1})f(u^*_{i_1})f(h^*_{j_1})\\ &\sum_{i_2}^{n_u} w_{ui_2}\sum_{j_2}^{n_h} w_{hi_2} f(\mathbf{y}_2|u^*_{i_2},h^*_{j_2})f(u^*_{i_2}|u^*_{i_1})f(h^*_{j_2}|h^*_{j_1})\dots\\ &\dots\sum_{i_T}^{n_u} w_{ui_T}\sum_{j_T}^{n_h} w_{hi_T} f(\mathbf{y}_T|u^*_{i_T},h^*_{j_T})f(u^*_{i_T}|u^*_{i_{T-1}})f(h^*_{j_T}|h^*_{j_{T-1}})\end{aligned} \tag{6}$$

where $\{u^*_i\}$, with $i=1,\dots,n_u$, and $\{h^*_j\}$, with $j=1,\dots,n_h$, are sets of Gauss-Legendre quadrature points, w_{ui} and w_{hj} are the corresponding weights and $[a,b]$ and $[d,e]$ are finite integration limits which replace the infinite ones for the random effects and the volatility process respectively. The choice of the grids and the number of evaluation points is crucial for the numerical precision. First, as proposed in Fridman and Harris (1998), the grids for the two latent processes are centered on $\mu_u = 0$ and $\mu_h = \alpha/(1-\delta)$ with a width of $3\sigma_u = 3\sigma_\eta/(\sqrt{1-\rho^2})$ and $3\sigma_h = 3\sigma_\nu/(\sqrt{1-\delta^2})$; this allows the grids to cover the support of the unconditional distributions with non negligible mass. Second, the number of quadrature points, n_u and n_h, are chosen according to the degree of smoothness of the integrands, i.e. the average distance between two points is less or equal to $\sigma_\eta/2$ for the random effect process and $\sigma_\nu/2$ for the volatility process.

4.2 Recursive algorithm

In order to facilitate the computation of the approximated likelihood (6) we use the matrix notation for the recursive algorithm. Define the $(n_u \times 1)$ and $(n_h \times 1)$

vectors of the initial probabilities as

$$\mathbf{f}_{u_1} = (f(u_1^*), f(u_2^*), \ldots, f(u_{n_u}^*))' \quad \text{and} \quad \mathbf{f}_{h_1} = (f(h_1^*), f(h_2^*), \ldots, f(h_{n_h}^*))'$$

and the $(n_u n_h \times 1)$ vector of the initial conditional probabilities as

$$\mathbf{f}_1 = (f(\mathbf{y}_1|u_1^*, h_1^*), f(\mathbf{y}_1|u_2^*, h_1^*), \ldots, f(\mathbf{y}_1|u_{n_u}^*, h_1^*), f(\mathbf{y}_1|u_1^*, h_2^*), \ldots, f(\mathbf{y}_1|u_{n_u}^*, h_{n_h}^*))'$$

In the same way, define the probability transition matrices of the two discretized latent variables as

$$\mathbf{F}_h = \begin{bmatrix} f(h_1^*|h_1^*) & \ldots & f(h_1^*|h_{n_h}^*) \\ \vdots & \vdots & \vdots \\ f(h_{n_h}^*|h_1^*) & \ldots & f(h_{n_h}^*|h_{n_h}^*) \end{bmatrix} \quad \mathbf{F}_u = \begin{bmatrix} f(u_1^*|u_1^*) & \ldots & f(u_1^*|u_{n_u}^*) \\ \vdots & \vdots & \vdots \\ f(u_{n_u}^*|u_1^*) & \ldots & f(u_{n_u}^*|u_{n_u}^*) \end{bmatrix}$$

and the $(n_u n_h \times 1)$ vector of conditional probabilities for $t = 2, \ldots, T$ as

$$\mathbf{f}_t = (\mathbf{f}_{t_1}, \mathbf{f}_{t_2} \ldots \mathbf{f}_{t_n_h})' \quad \text{with} \quad \mathbf{f}_{t_j} = (f(\mathbf{y}_t|u_1^*, h_j^*), f(\mathbf{y}_t|u_2^*, h_j^*), \ldots, f(\mathbf{y}_t|u_{n_u}^*, h_j^*))'$$

Finally, define the vectors $\mathbf{w}_u = (w_{u1}, w_{u2}, \ldots, w_{un_u})'$, $\mathbf{w}_h = (w_{h1}, w_{h2}, \ldots, w_{hn_h})'$ and the matrices $\mathbf{W}_u = (\mathbf{w}_u, \mathbf{w}_u, \ldots, \mathbf{w}_u)$, $\mathbf{W}_h = (\mathbf{w}_h, \mathbf{w}_h, \ldots, \mathbf{w}_h)$ of dimensions $(n_u \times n_u)$ and $(n_h \times n_h)$ respectively. The approximated marginal likelihood can be computed through the following iterative algorithm

1. compute the joint probability at the first time point $(t = 1)$:

$$\mathbf{l}_1 = \left(\frac{b-a}{2}\right)\left(\frac{e-d}{2}\right) \text{diag}(\mathbf{f}_1)((\mathbf{f}_{h_1} \circ \mathbf{w}_h) \otimes (\mathbf{f}_{u_1} \circ \mathbf{w}_u)) \tag{7}$$

   where $\circ$ indicates the Hadamard product between two matrices and $\text{diag}(\mathbf{f}_1)$ is the block-diagonal matrix created from the vector $\mathbf{f}_1$. The generic element of $\mathbf{l}_1$ is $l_{1ij} = f(\mathbf{y}_1, u_t = u_i^*, h_t = h_j^*)$.

2. compute the forward intermediate joint probabilities for $t = 2, \ldots, T$:

$$\mathbf{l}_t = \left(\frac{b-a}{2}\right)\left(\frac{e-d}{2}\right) \text{diag}(\mathbf{f}_t)((\mathbf{F}_h \circ \mathbf{W}_h) \otimes (\mathbf{F}_u \circ \mathbf{W}_u))\mathbf{l}_{t-1}, \tag{8}$$

   where $\text{diag}(\mathbf{f}_t)$ is the block-diagonal matrix created from the vector $\mathbf{f}_t$. The generic element of $\mathbf{l}_t$ is $l_{tij} = f(\mathbf{y}_1, \ldots, \mathbf{y}_t, u_t = u_i^*, h_t = h_j^*)$.

3. compute the approximated marginal likelihood (6) as

$$\tilde{L}(\boldsymbol{\theta}|\mathbf{y}) = \mathbf{1}'\mathbf{l}_T \tag{9}$$

   where $\mathbf{1}$ is a $(n_u n_h \times 1)$ vector of ones.

The maximization of expression (9) is carried out by using a quasi-Newton algorithm in which the gradient and the Hessian matrix are obtained through numerical derivatives. The choice of the starting values for $\boldsymbol{\theta}$ requires two

separate strategies, one for the parameters of the stochastic volatility component and one for the remaining parameters. As for the former, we adapt to repeated cross-sectional data the method used by Bartolucci and De Luca (2001) in a time series context. First, consider the stochastic volatility model for the error terms of the ARE model

$$y_{it}^* = y_{it} - \beta_0 - u_t - \mathbf{x}_{it}'\boldsymbol{\beta} = \exp(h_t/2)\epsilon_{it}, \tag{10}$$

from which, as observed by Harvey et al. (1994), the following expression can be derived

$$\log y_{it}^{*2} = h_t + \log \epsilon_{it}^2$$

Since $\mathrm{E}(\log \epsilon_{it}^2) = -1.27$ we have that

$$h_t \approx h_t^* = \log \hat{y}_{it}^{*2} + 1.27$$

where $\hat{y}_{it}$ are obtained as the level-1 residuals of the ARE model. Second, smooth h_t^* through a k-term moving average, run a regression of h_t^* on h_{t-1}^* and use the resulting estimates as initial values of the stochastic volatility parameters. As concerns the remaining parameters, the estimates of the fitted ARE model are taken as starting values.

4.3 Filtering, Smoothing and Prediction

In order to obtain optimal estimators of the unobserved-state vectors $\mathbf{u}$ and $\mathbf{h}$, we perform filtering and smoothing, which differ in the conditioning information set. Moreover, we derive the one-step-ahead predictors.

4.3.1 Filtering

Denoting with $\mathbf{Y}_t = (\mathbf{y}_1, \ldots, \mathbf{y}_t)$, the filtered volatility and the filtered random effects at time t are defined as

$$\hat{h}_t(\mathbf{Y}_t) = E(h_t|\mathbf{Y}_t) = \int h_t f(h_t|\mathbf{Y}_t) dh_t$$
$$\hat{u}_t(\mathbf{Y}_t) = E(u_t|\mathbf{Y}_t) = \int u_t f(u_t|\mathbf{Y}_t) du_t.$$

The filtered values are computed directly from the quantities obtained in the forward recursive algorithm:

$$\hat{h}_t(\mathbf{Y}_t) = \int h_t \frac{f(h_t, \mathbf{Y}_t)}{f(\mathbf{Y}_t)} dh_t = \frac{\int\int h_t f(\mathbf{Y}_t, u_t, h_t) du_t dh_t}{\int\int f(\mathbf{Y}_t, u_t, h_t) du_t dh_t} \simeq \frac{\sum_j^{n_h} h_j^* \sum_i^{n_u} l_{tij}}{\sum_i^{n_u}\sum_j^{n_h} l_{tij}}$$
$$\hat{u}_t(\mathbf{Y}_t) = \int u_t \frac{f(u_t, \mathbf{Y}_t)}{f(\mathbf{Y}_t)} du_t = \frac{\int\int u_t f(\mathbf{Y}_t, u_t, h_t) dh_t du_t}{\int\int f(\mathbf{Y}_t, h_t, u_t) du_t dh_t} \simeq \frac{\sum_i^{n_u} u_j^* \sum_j^{n_h} l_{tij}}{\sum_i^{n_u}\sum_j^{n_h} l_{tij}}$$

4.3.2 Smoothing

In the multilevel framework, the random effects are predicted by using either the Best Linear Unbiased Predictor or the Empirical Bayes Predictor. In both cases, the best predictor is the expected value conditioned to the whole observed sample or empirical posterior distribution (for an overview see Searle et al. (1992) and Skrondal and Rabe-Hesketh (2004)). In stochastic volatility models the smoothed values are:

$$\hat{h}_t(\mathbf{Y}_T) = E(h_t|\mathbf{Y}_T) = \int\int h_t \frac{f(\mathbf{Y}_t, h_t, u_t) f(\mathbf{y}_{t+1}, \ldots, \mathbf{y}_T|h_t, u_t)}{f(\mathbf{Y}_T)} du_t dh_t \quad (11)$$

$$\hat{u}_t(\mathbf{Y}_T) = E(u_t|\mathbf{Y}_T) = \int\int u_t \frac{f(\mathbf{Y}_t, h_t, u_t) f(\mathbf{y}_{t+1}, \ldots, \mathbf{y}_T|h_t, u_t)}{f(\mathbf{Y}_T)} dh_t du_t. \quad (12)$$

In both expressions, the density $f(\mathbf{Y}_T)$ is approximated through the marginal likelihood $\tilde{L}(\boldsymbol{\theta}|\mathbf{y})$ of Eq. (9) and $f(\mathbf{Y}_t, h_t, u_t)$ through l_{tij}. The intermediate conditional probabilities, $f(\mathbf{y}_{t+1}, \ldots, \mathbf{y}_T|h_t, u_t)$ can be obtained by using the following backward recursion

1. Define $\mathbf{b}_T = \mathbf{1}_{n_h n_u}$.
2. Compute the backward conditional probability at time t, with $t = T - 1, \ldots, 1$:

$$\mathbf{b}_t = \left(\frac{b-a}{2}\right)\left(\frac{e-d}{2}\right)\text{diag}(\mathbf{f}_{t+1})((\mathbf{F}_h \circ \mathbf{W}_h) \otimes (\mathbf{F}_u \circ \mathbf{W}_u))\mathbf{b}_{t+1}$$

   where the generic element is $b_{tij} = f(\mathbf{y}_{t+1}, \ldots, \mathbf{y}_T|u_t = u_i^*, h_t = h_j^*)$.

With this further recursion, the smoothed values are approximated as

$$\hat{h}_t(\mathbf{Y}_T) \simeq \frac{\sum_j^{n_h} h_j^* \sum_i^{n_u} l_{tij} b_{tij}}{\tilde{L}(\boldsymbol{\theta}|\mathbf{y})} \; ; \quad \hat{u}_t(\mathbf{Y}_T) \simeq \frac{\sum_i^{n_u} u_i^* \sum_j^{n_h} l_{tij} b_{tij}}{\tilde{L}(\boldsymbol{\theta}|\mathbf{y})}.$$

4.3.3 Prediction

The one-step-ahead predicted values of volatility and random effects are derived through the following approximations:

$$\begin{aligned}
\hat{h}_t(Y_{t-1}) = E(h_t|Y_{t-1}) &= \frac{\int h_t f(h_t|Y_{t-1}) f(Y_{t-1}) dh_t}{\int\int f(Y_{t-1}, u_{t-1}, h_{t-1}) dh_{t-1} du_{t-1}} \\
&= \frac{\int h_t \Big\{ \int f(h_t|h_{t-1}) \big[\int f(Y_{t-1}, u_{t-1}, h_{t-1}) du_{t-1} \big] dh_{t-1} \Big\} dh_t}{\int\int f(Y_{t-1}, u_{t-1}, h_{t-1}) du_{t-1} dh_{t-1}} \\
&\simeq \frac{\frac{e-d}{2} \sum_j^{n_h} h_j w_{hj} \sum_{j'}^{n_h} F_{hjj'} w_{hj'} \sum_i^{n_u} l_{t-1ij'}}{\sum_i^{n_u} \sum_j^{n_h} l_{t-1,ij}} \\
\hat{u}_t(Y_{t-1}) &\simeq \frac{\frac{b-a}{2} \sum_i^{n_u} u_i w_{ui} \sum_{i'}^{n_u} F_{uii'} w_{ui'} \sum_j^{n_h} l_{t-1,i'j}}{\sum_i^{n_u} \sum_j^{n_h} l_{t-1,ij}}
\end{aligned}$$

where $F_{uii'}$ and $F_{hjj'}$ denote the generic elements of the transition matrices F_u and F_h respectively.

5 Application to tribal Art prices

In this section we illustrate the application of our model to the first database of ethnic artworks. The responses are the logged prices for 28 semesters for the overall sample size of 13955 items. We take the fixed effects hedonic specification (FE) as the benchmark model. The covariates are selected through stepwise (forward and backward) techniques and then combining parsimony with Art Economics arguments. Table 2 reports the parameter estimates for the three models: fixed effects (FE), autoregressive random effects (ARE) and stochastic volatility with autoregressive random effects (SVARE), fitted on the same data set with the same set of covariates. The asymptotic standard errors for the FE and SVARE models are derived from the Hessian matrix of the likelihood functions. The robust standard error for the ARE model are derived by means of the wild bootstrap for multilevel models introduced in Modugno and Giannerini (2015).

In order to assess whether the specifications proposed manage to model satisfactorily the time dynamics and the heterogeneity observed, we have implemented a series of diagnostic tests. Table 3 reports information on the goodness of fit of the models whereas in Table 4 we present the results of some diagnostic tests and indicators on the residuals. In particular, the first row of Table 4 shows the p-values for the Shapiro-Wilk test for normality of the residuals of the three models; the `shapiro.test` function in `R` limits the sample size to 5000. The results presented are the median p-values over 20000 random subsamples of size 5000 drawn from the original sample. The last two rows show the indexes of skewness and kurtosis b_1 and b_2 as in Joanes and Gill (1998) computed on level-1 residuals. We derive the standardized residuals for the ARE model from the BLUP of the random effects whereas for the SVARE model we use the smoothed values for both the random effects (Eq. 12) and the volatility (Eq. 11).

As concerns the time dynamics we assess the adequateness of the models by computing the sample global and partial autocorrelation functions over time varying quantities such as level-2 residuals. Moreover, we use the metric entropy measure S_k defined as

$$S_k = \frac{1}{2} \int \int \left[\{ f_{(X_t, X_{t+k})}(x_1, x_2) \}^{1/2} - \{ f_{X_t}(x_1) f_{X_{t+k}}(x_2) \}^{1/2} \right]^2 dx_1 dx_2 \quad (13)$$

where f_{X_t} and $f_{(X_t, X_{t+k})}$ denote the probability density function of X_t and of the vector (X_t, X_{t+k}) respectively. The measure is a particular member of the family of relative entropies, which includes as a special case non-metric entropies often referred to as Shannon or Kullback–Leibler divergence. It can be interpreted as a nonlinear autocorrelation function and possesses many desirable properties. We use S_k as in Giannerini et al. (2015) to test for nonlinear serial dependence and as in Granger et al. (2004) to test for serial independence (see the supplementary material for more details). The tests are implemented in the `R` package `tseriesEntropy` Giannerini (2015). In the spirit of time series analysis, if the specification is appropriate then the residuals behave as a white noise process and diagnostic tests can suggest directions to improve the existing

model.

Finally, Table 5 summarizes and compares the prediction/forecasting capability of the three models under scrutiny. The aggregate measures of prediction error are the Mean Absolute (Prediction) Error (MAE) and the Root Mean Square (Prediction) Error (RMSE):

$$\text{MAE} = \frac{1}{n_{T+1}} \sum_{i=1}^{n_t} |y_{i,T+1} - \hat{y}_{i,T+1}| \; ; \; \text{RMSE} = \sqrt{\frac{1}{n_{T+1}} \sum_{i=1}^{n_{T+1}} \left(y_{i,T+1} - \hat{y}_{i,T+1}\right)^2} \tag{14}$$

The first two rows of Table 5 report the prediction error over 100 (out of sample) items within the time span 1998-2011, and the last two rows of the table show the forecasting performance over all the 73 observations of semester 2011-2. Such observations have not been included in the model so that the measures reflect a genuine forecasting performance.

5.1 FE and ARE models

The parameter estimates of the FE and ARE models are very similar (first two columns of Table 2), still, there are important differences: first, the ARE fit is more parsimonious and results in a smaller BIC and AIC (see Table 3); also, it provides a decomposition of the total variability of the response in between-time and within-time variability. Furthermore, in the FE model the time dynamics is modelled through 28 dummy variables while in the ARE model the time dynamics is fully captured through the AR(1) specification with the two parameters, ρ and σ_η, see Figure 1 and 2 of the Supplementary Material (SM from now on). The same tests performed on level-2 residuals of the ARE model show no structure (see Figure 3 and 4 of the SM). Finally, the ARE model provides a superior one-step-ahead forecasting of the price whereas the prediction performance is the same as that of the FE model (see Table 5). The Shapiro-Wilk test (see Table 4) points to a deviation from normality in level-1 residuals of the ARE model (whereas it does not reject the assumption of normality for level-2 residuals). As discussed above, this is consistent with the findings in literature and might be due to heteroscedasticity. In fact, similarly to other assets, level-1 residuals show a leptokurtic behaviour as shown in Figure 4 and by looking at the kurtosis index in Table 4. Furthermore, we reject the assumption of homogeneity of the variance across time points, tested through a non-parametric version of the Levene (1960) rank-based test (Kruskal and Wallis, 1952).

The plot of s_t^{ARE}, the standard deviations of level-1 residuals $\hat{\epsilon}_{it}$ in Figure 2(left) provides a visual evidence of volatility patterns. The entropy measure S_k shown in Figure 2(right) confirms the presence of a linear serial dependence (the test for nonlinearity does not reject, see Figure 6 of the supplementary material) and the correlograms of Figure 3 indicate a AR(1)-type dependence structure for the volatility. In the following subsection we account for the observed heterogeneity by fitting the multilevel model with autoregressive random effects and stochastic volatility (SVARE).

Table 2: Parameter estimates for models FE, ARE and SVARE with standard errors in parentheses. For each categorical variable the baseline category is indicated. The complete set of estimates is available in the supplementary material

	FE	ARE	SVARE
σ^2	0.225 (0.006)	0.226 (0.004)	-
σ^2_η	-	0.022 (0.015)	0.021 (0.000)
ρ	-	0.837 (0.140)	0.848 (0.025)
α	-	-	-0.142 (0.012)
δ	-	-	0.931 (0.035)
σ^2_ν	-	-	0.158 (0.001)
CABS (Yes vs No)	0.103 (0.013)	0.102 (0.015)	0.093 (0.012)
CABC (Yes vs No)	0.138 (0.010)	0.138 (0.011)	0.127 (0.010)
CAES (Yes vs No)	0.078 (0.015)	0.078 (0.016)	0.081 (0.013)
Christie's (vs Sotheby's)	-0.109 (0.011)	-0.110 (0.013)	-0.153 (0.011)
Paris (vs New York)	-0.126 (0.013)	-0.124 (0.015)	-0.076 (0.012)
Illustration: baseline Absent			
- Miscellaneous col.	0.452 (0.021)	0.453 (0.031)	0.441 (0.020)
- Col. cover	1.494 (0.076)	1.494 (0.104)	1.395 (0.075)
- Col. half page	0.996 (0.025)	0.995 (0.034)	0.864 (0.024)
- Col. full page	1.148 (0.026)	1.149 (0.034)	1.030 (0.026)
- More than one col.	1.376 (0.028)	1.377 (0.037)	1.273 (0.027)
- Col. quarter page	0.823 (0.021)	0.823 (0.029)	0.701 (0.020)
- Miscellaneous b/w	0.500 (0.047)	0.500 (0.040)	0.411 (0.038)
- b/w half page	0.647 (0.061)	0.647 (0.057)	0.553 (0.047)
- b/w full page	1.024 (0.239)	1.025 (0.290)	0.829 (0.190)
- b/w quarter page	0.383 (0.029)	0.383 (0.033)	0.302 (0.025)
Description on the catalogue: baseline Absent			
- Short visual descr.	-0.099 (0.037)	-0.100 (0.042)	-0.138 (0.035)
- Visual descr.	0.078 (0.038)	0.076 (0.044)	0.036 (0.035)
- Broad visual descr.	0.325 (0.041)	0.325 (0.047)	0.274 (0.038)
- Critical descr.	0.319 (0.041)	0.319 (0.049)	0.266 (0.038)
- Broad critical descr.	0.719 (0.045)	0.718 (0.055)	0.640 (0.042)
Type of object: baseline Furniture			
- Masks	0.121 (0.025)	0.121 (0.025)	0.111 (0.023)
- Religious objects	0.047 (0.027)	0.046 (0.030)	0.039 (0.025)
- Ornaments	-0.161 (0.026)	-0.161 (0.030)	-0.111 (0.024)
- Sculptures	0.062 (0.023)	0.061 (0.025)	0.062 (0.021)
- Musical instruments	-0.098 (0.036)	-0.098 (0.044)	-0.061 (0.033)
- Tools	-0.088 (0.024)	-0.089 (0.027)	-0.075 (0.022)
- Clothing	-0.077 (0.037)	-0.078 (0.040)	-0.075 (0.032)
- Weapons	-0.131 (0.030)	-0.131 (0.033)	-0.095 (0.028)
Continent: baseline Africa			
- America	0.175 (0.013)	0.174 (0.016)	0.162 (0.013)
- Oceania	0.167 (0.012)	0.167 (0.014)	0.147 (0.011)
Patina: baseline Not indicated			
- Pejorative	0.210 (0.046)	0.209 (0.044)	0.167 (0.046)
- Appreciative	0.124 (0.012)	0.124 (0.014)	0.104 (0.012)
Historicization: baseline Absent			
- Museum certification	0.045 (0.017)	0.045 (0.021)	0.021 (0.016)
- Relevant museum certification	0.065 (0.016)	0.064 (0.018)	0.045 (0.015)
- Simple certification	0.053 (0.011)	0.053 (0.013)	0.061 (0.010)

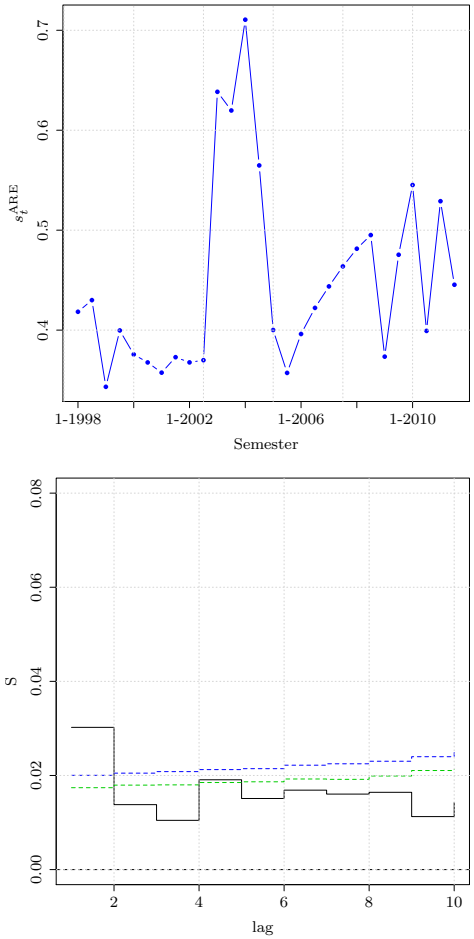

Figure 2: Standard deviations of ARE level-1 residuals s_t^{ARE}: time plot (left) and entropy measure of dependence (right). The confidence bands correspond to the null hypothesis of serial independence at levels 90% and 95% up to 10 lags/semesters.

Table 3: Loglikelihood, number of parameters and Information criteria for the hedonic regression (FE), ARE and SVARE models.

	FE	ARE	SVARE
loglik	-9405.833	-9472.96	-8785.624
n. par	70	45	47
AIC	18952	19036	17665
BIC	19480	19375	18020

Table 4: p-values of the Shapiro-Wilk test and indexes of skewness and kurtosis for the FE, ARE and SVARE (level-1) residuals.

	FE	ARE	SVARE
Shapiro-Wilk p-value	<2e-17	2e-16	3e-07
Skewness	-0.1683	-0.1727	0.0221
Kurtosis	1.0553	1.0630	0.4568

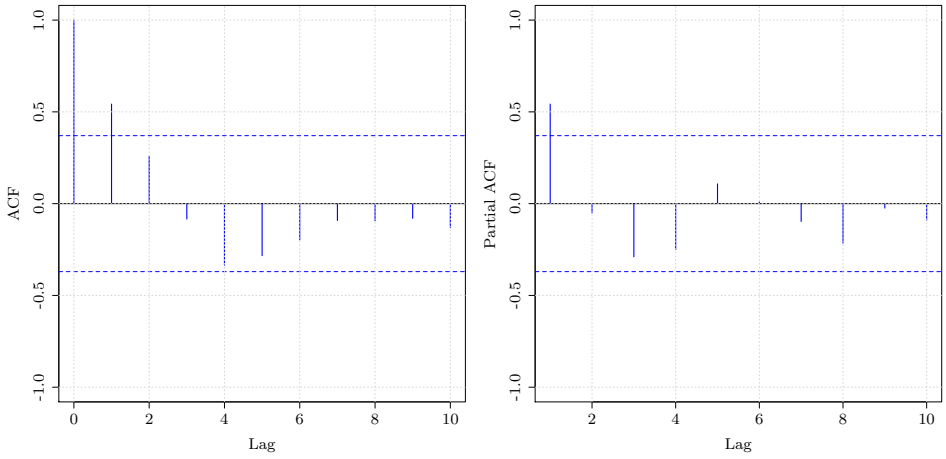

Figure 3: Global (left) and partial (right) empirical autocorrelation functions of standard deviations of ARE level-1 residuals.

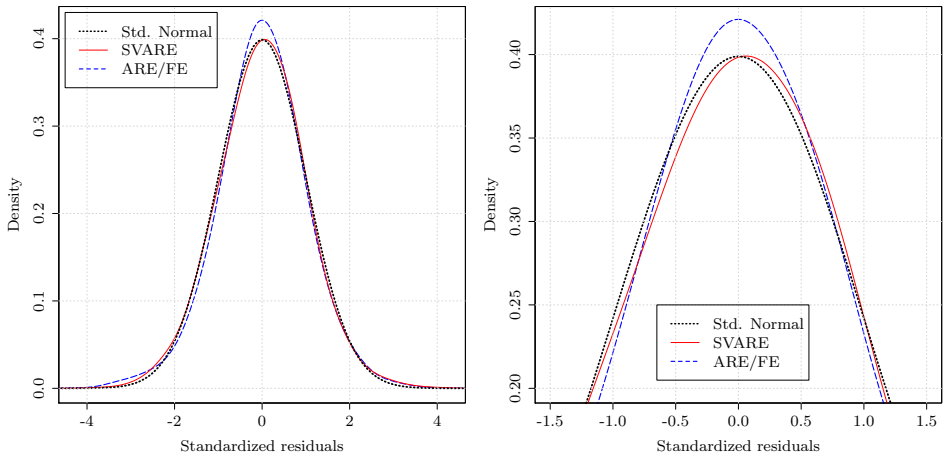

Figure 4: Density of level-1 residuals. The right panel is a zoomed detail of the peak that highlights the kurtosis.

5.2 SVARE model

A first important result is that the parameter estimates differ to some extent from those of the FE/ARE models, see Table 2. Even if in most cases the significance of the parameters does not change, SVARE estimates account properly for the volatility and reflect more closely the impact of the covariates on artwork prices. The estimate of the volatility parameter $\hat{\delta} = 0.931$ agrees with those of models for financial time series reported in literature and indicates a non-negligible volatility persistence. Indeed, the goodness of fit of the SVARE model increases considerably as witnessed by the information criteria shown in Table 3. Note also that the standard errors of the estimates are almost uniformly the smallest among the three models.

Also for the SVARE model, the Shapiro-Wilk test of normality of level-1 residuals rejects the null hypothesis (Table 4). Nevertheless, the leptokurtic behaviour of residuals is considerably reduced with respect to both the ARE and FE models. This is shown in Table 4 (the skewness disappears and the kurtosis is more than halved) and in Figure 4 where we show the densities of level-1 residuals for the ARE and SVARE models. Note the agreement of SVARE residuals with the standard Normal density (dotted in the figure). See also Figure 11 of the SM for a normal qqplot of level-1 residuals of the two models. As in the FE/ARE case, we compute the diagnostic tests of dependence on level-2 residuals $\hat{\eta}_t$ and $\hat{\nu}_t$. Both the correlograms and the entropy measure S_k indicate the absence of any dependence structure (see Figures 7 - 10 of the Supplementary Material) so that we may argue that the SVARE specification manages to capture the volatility dynamics. Finally, from Table 5 it emerges clearly that the SVARE specification performs best among competitors in terms of both prediction and forecasting.

Table 5: Prediction/forecasting performance of the three models over 100 out-of-sample units within the time span 1998-2011 (rows 1-2) and over 73 units of the out-of-sample semester, 2011-2 (rows 3-4).

		FE	ARE	SVARE
Prediction	MAE	0.288	0.288	0.278
	RMSE	0.355	0.355	0.344
Forecast	MAE	0.534	0.376	0.355
	RMSE	0.673	0.512	0.479

Note that the choice of the number of quadrature points n_u and n_h is crucial to the estimation of the SVARE model. We first set $n_u = 21$ and $n_h = 41$ according to the rule given in Section 4.1, then we increased them up to 71. As for the autoregressive part of the model, we found that the parameters' estimates are quite robust except for ρ whose estimate stabilizes for $n_u \geq 51$. The estimates of the Stochastic Volatility component are more sensitive to the number of quadrature points so that we set $n_u = n_h = 61$. This yields sensible results even for high values of δ, that is, in those cases where the integrand is very peaked and requires a large number of quadrature points to be well approximated.

6 Conclusions

A great added value of the SVARE specification for repeated cross sections is that it encompasses in a unique framework the trends in the mean and in the volatility of artwork prices. The left panel of Figure 5 compares the resulting biannual price indexes obtained through the ARE and SVARE fits. They are computed with fixed base in semester $b =$"1-1998" as

$$I_t = \frac{e^{\hat{\beta}_{0t}}}{e^{\hat{\beta}_{0b}}} \times 100,$$

where $\hat{\beta}_{0t}$ are the BLUP values $\hat{\beta}_0 + \hat{u}_t$. Note that the indexes are similar, nevertheless, the SVARE model provides the predicted volatility values $\hat{h}_t$ of Eq. 11 (right panel of Figure 5). This is important additional information on the predictability of the prices that complements that of the index and can be exploited by art market stakeholders for an informed decision making.

Acknowledgements

We would like to thank Guido Candela, Antonello Scorcu, Massimiliano Castellani and Pierpaolo Pattitoni for useful discussions. Silvia Cagnone and Lucia Modugno acknowledge the financial support from the grant RBFR12SHVV

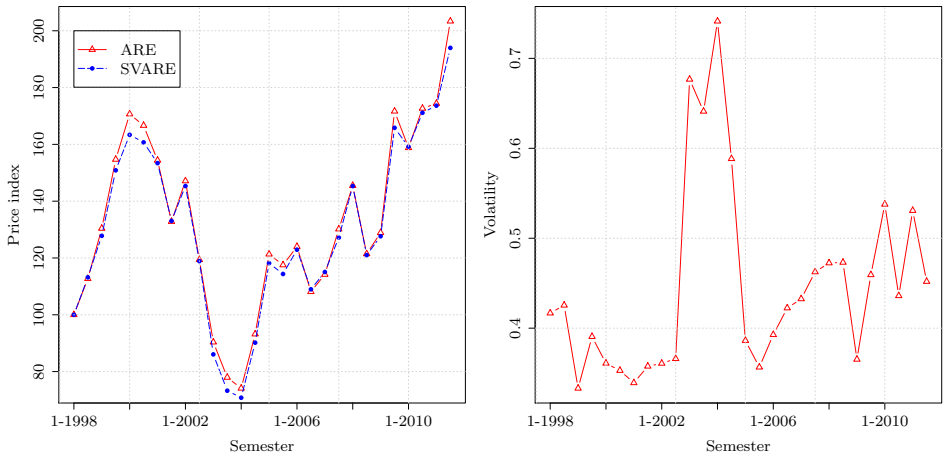

Figure 5: Price index (with base "semester 1-1998") of the art market for the ARE and SVARE models (left panel). Plot of volatility of the SVARE model (right panel).

funded by the Italian Government (FIRB project "Mixture and latent variable models for causal inference and analysis of socio-economic data").

References

Agnello, R. and R. Pierce (1996). Financial returns, price determinants, and genre effects in american art investment. *Journal of Cultural Economics 20*(4), 359–383.

Bartolucci, F. and G. De Luca (2001). Maximum likelihood estimation of a latent variable time-series model. *Applied Stochastic Models in Business and Industry 17*, 5–17.

Baumol, W. (1986). Unnatural value: or art investment as floating crap game. *The American Economic Review 76*(2), 10–14.

Bocart, F. and C. Hafner (2012). Econometric analysis of volatile art markets. *Computational Statistics & Data Analysis 56*(11), 3091 – 3104.

Bocart, F. Y. and C. M. Hafner (2013). Volatility of price indices for heterogeneous goods with applications to the fine art market. *Journal of Applied Econometrics*. doi: 10.1002/jae.2355.

Browne, W. and H. Goldstein (2010). MCMC sampling for a multilevel model with nonindependent residuals within and between cluster units. *Journal of Educational and Behavioral Statistics 35*, 453–473.

Cagnone, S. and F. Bartolucci (2016). Adaptive quadrature for maximum likelihood estimation of a class of dynamic latent variable models. *submitted*.

Candela, G., M. Castellani, and P. Pattitoni (2013). Reconsidering psychic return in art investments. *Economics Letters 118*(2), 351 – 354.

Chanel, O. (1995). Is art market behaviour predictable? *European Economic Review 39*, 519–527.

Chanel, O., L. A. Grard-Varet, and V. Ginsburgh (1996). The relevance of hedonic price indices. The case of paintings. *Journal of Cultural Economics 20*, 1–24.

Collins, A., A. E. Scorcu, and R. Zanola (2009). Reconsidering hedonic art price indexes. *Economics Letters 104*, 57–60.

DiPrete, T. A. and D. B. Grusky (1990). The multilevel analysis of trends with repeated cross-sectional data. *Sociological Methodology 20*, 337–368.

Fridman, M. and L. Harris (1998). A Maximum Likelihood Approach for Non-Gaussian Stochastic Volatility Models. *Journal of Business & Economic Statistics 16*(3), 284–291.

Giannerini, S. (2015). *`tseriesentropy`: Entropy Based Analysis and Tests for Time Series*. R package version 0.5-13.

Giannerini, S., E. Maasoumi, and E. Bee Dagum (2015). Entropy testing for nonlinear serial dependence in time series. *Biometrika 102*, 661–675.

Ginsburgh, V. and P. Jeanfils (1995). Long-term comovements in international markets for paintings. *European Economic Review 39*, 538–548.

Goetzmann, W., E. Mamonova, and C. Spaenjers (2014, August). The economics of aesthetics and three centuries of art price records. Working Paper 20440, National Bureau of Economic Research.

Goetzmann, W. N. (1993). Accounting for taste: Art and financial markets over three centuries. *The American Economic Review 83*, 1370–1376.

Goldstein, H. (2010). *Multilevel Statistical Models* (4th ed.). Chichester: J. Wiley & Sons.

Granger, C. W. J., E. Maasoumi, and J. Racine (2004). A dependence metric for possibly nonlinear processes. *J. Time Ser. Anal. 25*(5), 649–669.

Harvey, A., E. Ruiz, and N. Shephard (1994). Multivariate stochastic variance models. *The Review of Economic Studies 61*(2), pp. 247–264.

Hodgson, D. and K. Vorkink (2004). Asset pricing theory and the valuation of canadian paintings. *The Canadian Journal of Economics / Revue canadienne d'Economique 37*(3), pp. 629–655.

Joanes, D. and C. Gill (1998). Comparing measures of sample skewness and kurtosis. *Journal of the Royal Statistical Society. Series D (The Statistician) 47*(1), 183–189.

Kitagawa, G. (1987). Non-gaussian state-space modeling of nonstationary time series (with discussion). *Journal of the American Statistical Association 82*, 1032–1063.

Kruskal, W. H. and W. A. Wallis (1952). Use of ranks in one-criterion variance analysis. *Journal of the American Statistical Association 47*, 583–621.

Levene, H. (1960). Robust tests for equality of variances. In I. Olkin (Ed.), *Contributions to Probability and Statistics*. Palo Alto, CA: Stanford Univ. Press.

Locatelli Biey, M. and R. Zanola (2005). The market for picasso prints: A hybrid model approach. *Journal of Cultural Economics 29*, 127–136.

Modugno, L., S. Cagnone, and S. Giannerini (2015). A multilevel model with time series components for the analysis of tribal art prices. *Journal of Applied Statistics 42*(10), 2141–2158.

Modugno, L. and S. Giannerini (2015). The wild bootstrap for multilevel models. *Communications in Statistics – Theory and Methods 44*(22), 4812–4825.

Rosen, S. (1974). Hedonic prices and implicit markets: product differentiation in pure competition. *Journal of Political Economy 82*, 34–55.

Searle, S. R., G. Casella, and C. E. McCulloch (1992). *Variance components*. New York: Wiley.

Skrondal, A. and S. Rabe-Hesketh (2004). *Generalized latent variable modeling: multilevel, longitudinal, and structural equation models*. New York: Chapman & Hall/CRC.

Tanizaki, H. and R. Mariano (1998). Nonlinear and nonnormal state-space modeling with monte-carlo stochastic simulations. *Journal of Econometrics 83*, 263–290.

Wooldridge, J. M. (2002). *Econometric Analysis of Cross Section and Panel Data*. Cambridge, Massachusetts: The MIT Press.